Cannon–Thurston maps, subgroup distortion, and hyperbolic hydra

Owen Baker*

Department of Mathematics & Statistics, McMaster University, 1280 Main Street West, Hamilton, Ontario, L8S 4K1, Canada

Timothy Riley†

Department of Mathematics, 310 Malott Hall, Cornell University, Ithaca, NY 14853, USA

September 6, 2012

Abstract

We prove that Cannon–Thurston maps are well–defined (that is, subgroup inclusion induces a map of the boundaries) for heavily distorted free subgroups inside the family of hyperbolic groups known as *hyperbolic hydra*. Whilst this indicates that distortion is not an obstacle to the map being well–defined, we show that heavy subgroup distortion always manifests in Cannon–Thurston maps (when they are well–defined) in that their continuity is quantifiably wild.

2010 Mathematics Subject Classification: 20F67

Key words and phrases: Cannon–Thurston map, hyperbolic group, subgroup distortion, hydra, Ackermann's function

1 Introduction

A *Cannon–Thurston* map $\partial\Lambda \to \partial\Gamma$ is a map between Gromov–boundaries that is induced by the inclusion $\Lambda \hookrightarrow \Gamma$ of an infinite hyperbolic subgroup Λ in a hyperbolic group Γ. (More details are in Section 2.)

For a finitely generated subgroup Λ of a finitely generated group Γ, define the *distortion function*

$$\mathrm{Dist}_\Lambda^\Gamma(n) := \max\{\, d_\Lambda(e,h) \mid h \in \Gamma,\ d_\Gamma(e,h) \le n \,\},$$

where d_Γ and d_Λ are word metrics with respect to some finite generating sets. We say that $f \preceq g$ for $f,g : \mathbb{N} \to \mathbb{N}$ when there exists $C>0$ such that $f(n) \le Cg(Cn+C)+Cn+C$ for all $n \ge 0$. We say $f \simeq g$ when $f \preceq g$ and $g \preceq f$. Up to $\simeq$, $\mathrm{Dist}_\Lambda^\Gamma(n)$ does not depend on the choices of finite generating sets.

If Λ is an undistorted subgroup in a hyperbolic group Γ (that is, $\mathrm{Dist}_\Lambda^\Gamma(n) \preceq n$), then Λ is also hyperbolic (e.g. [4, page 461]) and the induced map $\partial\Lambda \to \partial\Gamma$ is readily seen to be well–defined.

*obaker@math.mcmaster.ca

†tim.riley@math.cornell.edu. Partial support from NSF grant DMS–1101651 and from Simons Foundation Collaboration Grant 208567 is gratefully acknowledged.

Generalizing an example of Cannon & Thurston [7], Mitra showed it also is well–defined when Λ is an infinite normal subgroup of Γ [16], and when Γ is a finite graph of groups where the edge inclusions are quasi–isometric embeddings and Λ is one of the (infinite) vertex– or edge–groups [17]. We recently gave the first example where the Cannon–Thurston map is known not to be well–defined [1], and Matsuda & Oguni showed it leads to examples where the subgroup in question can be any non–elementary hyperbolic group or, even, relatively hyperbolic group [15].

In this article we show that extreme distortion is no barrier to the Cannon–Thurston map being well–defined, but such distortion manifests in extreme wildness of the map.

The *hyperbolic hydra* Γ_k $(k = 1, 2, \ldots)$ of [3] are a family of hyperbolic groups with finite–rank free subgroups Λ_k exhibiting the fastest–growing distortion functions known for hyperbolic subgroups of hyperbolic groups: $\mathrm{Dist}_{\Lambda_k}^{\Gamma_k}$ grows at least like A_k, the k–th of Ackermann's family of increasingly fast–growing functions that begins with $A_1(n) = 2n$, $A_2(n) = 2^n$, and $A_3(n)$ equalling the height–n tower of powers of 2. We prove:

Theorem 1.1. *Hyperbolic hydra have Cannon–Thurston maps $\partial\Lambda_k \to \partial\Gamma_k$ for all k.*

Wildness is apparent in Cannon & Thurston's example [7], a closed hyperbolic 3–manifold M fibering over the circle with fiber a hyperbolic surface S and pseudo–Anosov monodromy. They show that the Cannon–Thurston map $\partial(\pi_1 S) = S^1 \to S^2 = \partial(\pi_1 M)$ for $\pi_1 S \hookrightarrow \pi_1 M$ is well–defined but is a group–equivariant space–filling curve.

We detect wildness in Cannon–Thurston maps in the relationship between "ε" and "δ" in their continuity (having given the boundaries their natural "*visual*" metrics—see Section 2).

The quality of a function $f : U \to V$ between metric spaces is measured by its *modulus of continuity* $\varepsilon : [0,\infty) \to [0,\infty]$,

$$\varepsilon(\delta) := \sup\{d_V(f(a), f(b)) \mid a, b \in U \text{ with } d_U(a,b) \leq \delta\},$$

a notion going back at least to Lebesgue in 1909 [14]. An upper bound on $\varepsilon(\delta)$ expresses a degree of good behaviour: f is uniformly continuous if $\varepsilon(\delta) \to 0$ as $\delta \to 0$; is Lipschitz if $\varepsilon(\delta) \leq C\delta$ for a constant $C > 0$; and is α–Hölder if $\varepsilon(\delta) \leq C\delta^\alpha$ for a constant $C > 0$. Wildness manifests in lower bounds on $\varepsilon(\delta)$, expressing that $\varepsilon(\delta)$ is extravagantly larger than δ when δ is close to 0. This happens for the Cannon–Thurston map (if it is well–defined) when the subgroup in question is heavily distorted:

Theorem 1.2. *Suppose $\Lambda \leq \Gamma$ are hyperbolic, Λ is non–elementary, the Cannon–Thurston map $\partial\Lambda \to \partial\Gamma$ is well–defined, and $r, s > 1$ are any respective visual parameters for visual metrics on $\partial\Lambda$ and $\partial\Gamma$. Then there exist $\alpha, \beta > 0$ so that for all $n \geq 0$ the modulus of continuity satisfies*

$$\varepsilon\left(\frac{\beta}{r^{\mathrm{Dist}_\Lambda^\Gamma(n)}}\right) \geq \frac{\alpha}{s^n}.$$

The reason we insist that Λ be non–elementary (i.e. contains an F_2 subgroup; see e.g. [11, Theorem 2.28]) in this theorem is that elementary Λ are not interesting in this context: if Λ is finite, then $\partial\Lambda$ is empty; if Λ is virtually $\mathbb{Z}$, then Λ is quasiconvex in G, its boundary $\partial\Lambda$ is two points, and the Cannon–Thurston map is well–defined and is an embedding.

The moduli of continuity of Cannon–Thurston maps have received attention before. Cannon–Thurston maps were originally defined from the Gromov–boundary of a Kleinian group to $\partial\mathbb{H}^3 = S^2$ — see [17]. (The Cannon–Thurston maps we study here for hyperbolic groups are a natural generalization of the cocompact case.) When the Kleinian group is a finitely generated Fuchsian group of the first type with bounded geometry and no parabolic elements, Miyachi [18] gives an upper bound on the modulus of continuity for the Cannon–Thurston map, and shows it is not Hölder continuous when the group is geometrically infinite.

Theorems 1.1 and 1.2 combine with [3, Theorem 1.1], which says that $\mathrm{Dist}_{\Lambda_k}^{\Gamma_k} \succeq A_k$, to give the following corollary which we will make precise and prove in Section 5.

Corollary 1.3. *For all $k \geq 2$, the modulus of continuity $\varepsilon(\delta)$ for the Cannon–Thurston map $\partial\Lambda_k \to \partial\Gamma_k$ for hyperbolic hydra grows at least like $1/n$ when δ grows like $1/A_k(n)$.*

By Lemma 3.2, this result neither depends on choices of finite generating sets for Γ_k and Λ_k, nor of visual metrics on $\partial\Gamma_k$ and $\partial\Lambda_k$ (but the constants involved may depend on such choices).

Remark 1.4. A detailed understanding of the Cannon–Thurston Map $\partial\Lambda_k \to \partial\Gamma_k$ appears hard to obtain, since, whilst $\partial\Lambda_k$ is a Cantor set (as Λ_k is free), $\partial\Gamma_k$ is not so readily identified. (I. Kapovich & M. Lustig [12] recently made advances in the understanding of Cannon–Thurston maps for certain free–by–cyclic groups, but $\Lambda_k \leq \Gamma_k$ do not fall within the scope of their work.) Here is what we can say about $\partial\Gamma_k$.

Splittings of hyperbolic free–by–cyclic groups $F \rtimes_\varphi \mathbb{Z}$ are studied in [13] and [5], the former dealing with the case where φ is an irreducible hyperbolic free group automorphisms, and the latter with φ a general hyperbolic free group automorphism. The argument preceding Corollary 15 in [13] shows that any hyperbolic free–by–cyclic group has one–dimensional boundary: the cohomological dimension of any (finitely generated free)–by–cyclic group is 2 (see e.g. [6, pp.185–7]), so [2, Corollary 1.4(d)] implies $\partial\Gamma_k$ has dimension $2-1=1$.

The argument of [13, Corollary 15] shows that any hyperbolic free–by–cyclic group has connected, locally connected boundary without global cut points. To see this, it suffices by [11, Theorems 7.1 and 7.2] to check that $F \rtimes_\varphi \mathbb{Z}$ is freely indecomposable, which is true for *any* free group automorphism φ. Indeed, the Bass–Serre tree T for any graph of groups decomposition of $F \rtimes_\varphi \mathbb{Z}$ admits a minimal action by the normal subgroup F with quotient a finite graph by Grushko's Theorem. This shows the edge stabilizers for the action of $F \rtimes_\varphi \mathbb{Z}$ on T are non–trivial, so the decomposition cannot be free.

On the other hand, Γ_k splits as an HNN–extension over $\mathbb{Z}$ for every k, so [11, Theorem 7.2] implies $\partial\Gamma_k$ has local cut points. Indeed, Γ_1 splits over $\mathbb{Z}$ as $\langle B, a_1\rangle *_{\langle a_1^{-1}t^2ua_1 = t^2v^{-1}\rangle}$, where B denotes the subgroup generated by all the defining generators other than a_0 and a_1 (then a_0 appears as $t^{-1}a_1t$), and for $k \geq 2$, Γ_k splits as an HNN–extension over $\mathbb{Z}$ with the stable letter a_k conjugating t to ta_{k-1}^{-1}.

The organization of the remainder of this article. In Section 2 we give background on hyperbolic groups and their boundaries. In Section 3 we define Cannon–Thurston maps and prove an embellished version of a lemma of Mitra giving necessary and sufficient conditions for their being well–defined. In Section 4 we review hyperbolic hydra groups and prove Theorem 1.1. In Section 5 we prove Theorem 1.2 and Corollary 1.3.

Acknowledgment. We thank Mahan Mj (formerly Mitra) for conversations which fueled an interest in the relationship between subgroup distortion and Cannon–Thurston maps.

2 Hyperbolic groups and their boundaries

This section contains a brief account of some pertinent background. More general treatments can be found in, for example, [4, 9, 11, 19] and Gromov's foundational article [10].

For a metric space X, the *Gromov product* $(a \cdot b)_e$ (or $(a \cdot b)_e^X$ to avoid ambiguity) of $a, b \in X$ with respect to $e \in X$ is

$$(a \cdot b)_e \;=\; \frac{1}{2}(d(a,e) + d(b,e) - d(a,b)).$$

One says X is (δ)*–hyperbolic* when

$$(a \cdot b)_e \;\geq\; \min\{(a \cdot c)_e, (b \cdot c)_e\} - \delta$$

for all $e, a, b, c \in X$, and X is hyperbolic when it is (δ)–hyperbolic for some $\delta \geq 0$. When X is a geodesic space this is equivalent to other standard definitions of hyperbolicity (such as δ–thin or δ–slim triangles), although the δ involved may not agree.

When X is (0)–hyperbolic and geodesic—that is, an $\mathbb{R}$–tree—$(a\cdot b)_e$ is the distance from e to the geodesic between a and b. Correspondingly, in a (δ)–hyperbolic geodesic space every pair of geodesics from e to a and to b, both parametrized by arc–length, 6δ–fellow–travel for approximately $(a\cdot b)_e$ and then diverge (by the *insize* characterization of hyperbolicity of [4, page 408]). Indeed:

Lemma 2.1. *In a (δ)–hyperbolic geodesic metric space, for every geodesic $[a,b]$ connecting a and b*

$$|d(e,[a,b])-(a\cdot b)_e| \le 6\delta.$$

Proof. See [4]: the proof of Proposition 1.22 on page 411 shows that *insizes* of geodesic triangles are at most 6δ, and the proof of Proposition 1.17(3) $\Longrightarrow$ (2) on page 409 shows that all geodesic triangles are 6δ–*thin*, and the claimed inequality follows. □

The *(Gromov–) boundary* ∂X of a hyperbolic metric space X is defined with reference to, but is in fact independent of, a point $e\in X$. It is the set of equivalence classes of sequences (a_n) in X such that $(a_m\cdot a_n)_e\to\infty$ as $m,n\to\infty$, where two such sequences (a_n) and (b_n) are equivalent when $(a_m\cdot b_n)_e\to\infty$ as $m,n\to\infty$. Indeed, they are equivalent when $(a_n\cdot b_n)_e\to\infty$ as $n\to\infty$ since

$$(a_n\cdot b_m)_e \ge \min\{(a_n\cdot b_n)_e,(b_m\cdot b_n)_e\}-\delta$$

by (δ)–hyperbolicity. Denote the equivalence class of (a_n) by $\lim a_n$.

When X is a *geodesic* hyperbolic metric space, there are equivalent definitions of ∂X, such as ∂X is the set of equivalence classes of geodesic rays emanating from x, where two such rays are equivalent when they stay uniformly close. The condition $(a_m\cdot a_n)_e\to\infty$ is what makes a sequence (a_n) ray–like, and the condition $(a_m\cdot b_n)_e\to\infty$ corresponds to uniform closeness.

Extend the Gromov product to $\overline{X}:=X\cup\partial X$ by

$$(a\cdot b)_e = \sup\liminf_{m,n\to\infty}(a_m\cdot b_n)_e$$

where the sup is over all sequences (a_m) and (b_n) in X representing (when in ∂X) or tending to (when in X) a and b, respectively. (The "$\sup\liminf$" is necessary—see [4, page 432].)

We note, for (3) in the following lemma, that in a proper geodesic hyperbolic metric space X, each pair of distinct points $a,b\in\partial X$ is joined by a bi–infinite geodesic line $[a,b]$ (Lemma 3.2 on page 428 of [4]).

Lemma 2.2. *Suppose X is a proper geodesic (δ)-hyperbolic metric space.*

(0). *If $x,y\in\overline{X}$ and $e\in X$, then there exist sequences (x_n) and (y_n) in X with $x=\lim x_n$, $y=\lim y_n$, and $(x\cdot y)_e=\lim_n(x_n\cdot y_n)_e$.*

(1). *If $a,b,c\in\overline{X}$ and $e\in X$, then $(a\cdot b)_e \ge \min\{(a\cdot c)_e,(c\cdot b)_e\}-2\delta$.*

(2). *If $a,b\in X$ and $c\in\partial X$, then $|d(a,b)-(a\cdot c)_b-(b\cdot c)_a| \le \delta$.*

(3). *If $e\in X$ and $[a,b]$ is any geodesic joining any $a,b\in\partial X$, then $|d(e,[a,b])-(a\cdot b)_e| \le 8\delta$.*

Proof. (0) and (1) are parts 3 and 4 of [4, page 433, Remark 3.17]. (Alternatively, see parts 3 and 5 of [19, Lemma 4.6] but note that there $\inf\liminf$ is used in place of $\sup\liminf$ and so the constants differ.)

For (2), using (0) take sequences c_n,c'_n both approaching c such that $(a\cdot c)_b=\lim(a\cdot c_n)_b$ and $(b\cdot c)_a=\lim(b\cdot c'_n)_a$. Now, (δ)-hyperbolicity yields $(a\cdot c_n)_b\ge\min\{(a\cdot c'_n)_b,(c_n\cdot c'_n)_b\}-\delta$. As $n\to\infty$, we have $(c_n\cdot c'_n)_b\to\infty$, but $(a\cdot c'_n)_b$ is bounded above by $d(a,b)$. So $(a\cdot c'_n)_b\le(a\cdot c_n)_b+\delta$ for all sufficiently large n. Interchanging the roles of c_n and c'_n we find $|(a\cdot c_n)_b-(a\cdot c'_n)_b|\le\delta$ for all sufficiently large n. Hence:

$$|d(a,b)-(a\cdot c_n)_b-(b\cdot c'_n)_a| \le \delta+|d(a,b)-(a\cdot c'_n)_b-(b\cdot c'_n)_a| = \delta+|0| = \delta$$

for all sufficiently large n. Taking the limit as $n \to \infty$ gives the result.

For (3) (cf. Exercise 3.18(3) [4, page 433]), choose sequences $a_n \to a$ and $b_n \to b$ along $[a,b]$. Also choose $a'_n, b'_n \in X$ as in (0) so that $(a \cdot b)_e = \lim (a'_n \cdot b'_n)_e$. For large enough n, the closest point of $[a,b]$ to e lies on $[a_n, b_n]$, so

$$|d(e,[a,b]) - (a_n \cdot b_n)_e| \;=\; |d(e,[a_n,b_n]) - (a_n \cdot b_n)_e| \;\leq\; 6\delta, \tag{1}$$

with the inequality being by Lemma 2.1. By the (δ)–hyperbolicity condition,

$$\begin{aligned}(a_n \cdot b_n)_e \;&\geq\; \min\{(a_n \cdot a'_n)_e, (a'_n \cdot b'_n)_e, (b_n \cdot b'_n)_e\} - 2\delta, \text{ and}\\ (a'_n \cdot b'_n)_e \;&\geq\; \min\{(a_n \cdot a'_n)_e, (a_n \cdot b_n)_e, (b_n \cdot b'_n)_e\} - 2\delta.\end{aligned}$$

As $n \to \infty$ both $(a_n \cdot a'_n)_e \to \infty$ and $(b_n \cdot b'_n)_e \to \infty$, but $\limsup (a_n \cdot b_n)_e$ and $\limsup (a'_n \cdot b'_n)_e$ are bounded above by $(a \cdot b)_e + 1$ (else, passing to subsequences, we can assume $(a_n \cdot b_n)_e > (a \cdot b)_e + 1/2$ for all n, and so $\liminf (a_n \cdot b_n)_e \geq (a \cdot b)_e + 1/2$ contrary to the definition of $(a \cdot b)_e$). So these two inequalities together give $|(a_n \cdot b_n)_e - (a'_n \cdot b'_n)_e| \leq 2\delta$ for all sufficiently large n. Combining this with (1) gives the result. □

Visual metrics are natural metrics on the boundary ∂X of a (δ)–hyperbolic space X. Their essence is that $a, b \in \partial X$ are close when geodesics from a basepoint $e \in X$ to a and from e to b fellow travel for a long distance. One might hope that if $r > 1$, then $d(a,b) = r^{-(a \cdot b)_e}$ would define such a metric, but unfortunately, as such, transitivity can fail. Instead, say that a metric d on ∂X is a *visual metric* with *visual parameter* $r > 1$ when there exist $k_1, k_2 > 0$ such that for all $a, b \in \partial X$,

$$k_1 r^{-(a \cdot b)_x} \;\leq\; d(a,b) \;\leq\; k_2 r^{-(a \cdot b)_x}. \tag{2}$$

By [4, Proposition 3.21, page 435] (existence) and [11, Theorem 2.18] (Hölder–equivalence):

Lemma 2.3. *Suppose X is a (δ)–hyperbolic space, $r > 1$, and $e \in X$ is the base point with respect to which ∂X is defined. Then there is a visual metric on ∂X with parameter r. Moreover, any two visual metrics d and d' on ∂X (perhaps with different r and e) are Hölder–equivalent in that there exists $\alpha > 0$ such that the identity map $(\partial X, d) \to (\partial X, d')$ is α–Hölder and its inverse is $(1/\alpha)$–Hölder.*

We will need that $\overline{X} := X \cup \partial X$ is a compactification of X:

Lemma 2.4. *If X is a proper (δ)–hyperbolic geodesic metric space, then there is compact, sequentially compact metric on $\overline{X} := X \cup \partial X$ such that the inclusions of X and of ∂X are homeomorphic onto their images. Under this topology, ∂X is closed, and a sequence x_n in X converges to $x \in \partial X$ if and only if (x_n) is in the equivalence class of x.*

Proof. Sequential compactness is proved in [4, page 430, III.H.3.7] and compactness in [4, page 433 III.H.3.18(4)]. The last sentence follows from [4, page 431, III.H.3.13]. □

3 Cannon–Thurston maps and Mitra's Lemma

Consider hyperbolic metric spaces (X, d_X) and (Y, d_Y) and a map $f : Y \to X$. The induced map $\hat{f} : \partial Y \to \partial X$, if it is well–defined, sends the point in ∂Y represented by the sequence (a_n) to the point in ∂X represented by $(f(a_n))$. In the case where X is a hyperbolic group, Y is a hyperbolic subgroup, both with word metrics associated to finite generating sets, and f is the inclusion map, $\hat{f}$ is the *Cannon–Thurston map*.

The following lemma is an embellished version of Mitra's criterion for the Cannon–Thurston map to be well–defined ([16, Lemma 2.1] and [17, Lemma 2.1]). It also addresses continuity: if the map is well–defined, then it is necessarily continuous.

Our notation is that $B_X(e,R) = \{x \in X \mid d_X(e,x) < r\}$ and $\overline{B}_X(e,R) = \{x \in X \mid d_X(e,x) \leq r\}$, and we write $\gamma = [x,y]_X$ to mean γ *is a geodesic in X from x to y*. The metrics on ∂X and ∂Y implicit in this lemma are any visual metrics $d_{\partial X}$ and $d_{\partial Y}$.

The additional hypothesis for (e) can be removed if $\delta_Y = 0$; for $\delta_Y > 0$, we do not know whether (e) is equivalent to or strictly weaker than (a)–(d) in its absence. An inclusion map of a subgroup into an ambient group is Lipschitz when both have word metrics coming from some finite generating sets, and so it is satisfied in that setting.

Lemma 3.1. *Suppose (X, d_X) and (Y, d_Y) are infinite proper geodesic hyperbolic metric spaces, and $f : Y \to X$ is a proper map. Fix a basepoint $e \in Y$. Define $M, M', M'' : [0,\infty) \to [0,\infty)$ by*

$$\begin{aligned}
M(N) &:= \inf\left\{(f(x)\cdot f(y))^X_{f(e)} \mid x,y \in Y \text{ and } (x\cdot y)^Y_e \geq N\right\},\\
M'(N) &:= \inf\left\{d_X(f(e),\gamma) \mid \gamma = [f(x), f(y)]_X \text{ for some } [x,y]_Y \text{ in } Y \smallsetminus B_Y(e,N)\right\},\\
M''(N) &:= \inf\left\{d_X(f(e),\gamma) \mid \gamma = [f(z), f(y)]_X \text{ for some } z \text{ on some } [e,y]_Y \text{ with } d_Y(e,z) \geq N\right\}.
\end{aligned}$$

The following are equivalent:

(a). *The induced map $\hat{f} : \partial Y \to \partial X$ is well–defined.*

(b). *The induced map $\hat{f} : \partial Y \to \partial X$ is well–defined and is uniformly continuous.*

(c). *$M(N) \to \infty$ as $N \to \infty$.*

(d). *$M'(N) \to \infty$ as $N \to \infty$.*

Moreover, if $\sup\{d_X(f(x), f(y)) \mid d_Y(x,y) \leq r\} < \infty$ *for all* $r \geq 0$*, then these are also equivalent to*

(e). *$M''(N) \to \infty$ as $N \to \infty$.*

Proof. Each of $M(N)$, $M'(N)$, and $M''(N)$ is a non–decreasing function, so is either bounded of tends to ∞ as $N \to \infty$.

That (b) $\Longrightarrow$ (a) is immediate.

Let us prove (a) $\Longrightarrow$ (c). Suppose $M(N) \leq C$ for all N. So there are sequences (p_n) and (q_n) in Y with $(p_n \cdot q_n)^Y_e \to \infty$ but $(f(p_n)\cdot f(q_n))^X_{f(e)} \leq C$ for all n. As $Y \cup \partial Y$ is sequentially compact by Lemma 2.4, both (p_n) and (q_n) have convergent subsequences. But the condition $(p_n \cdot q_n)^Y_e \to \infty$ precludes any such subsequence from converging in Y, so those subsequences converge in ∂Y, indeed to the same point.

Next we prove (c) $\Longrightarrow$ (b). Suppose sequences (p_n) and (q_n) in Y both represent the same point in ∂Y. Then $(p_n \cdot q_n)^Y_e \to \infty$ as $n \to \infty$, and so $(f(p_n)\cdot f(q_n))^X_{f(e)} \to \infty$, since $M(N) \to \infty$ as $N \to \infty$. Thus if $(f(p_n))$ and $(f(q_n))$ converge to points in ∂X, then those points are the same.

So, to prove $\hat{f}$ is well–defined, it suffices to show that if a sequence (a_n) in Y represents a point in ∂Y (and so $d_Y(e,a_n) \to \infty$ by properness of Y), then $(f(a_n))$ represents a point in ∂X. Indeed, it suffices to show that a *subsequence* of $(f(a_n))$ represents a point in ∂X. By sequential compactness of $X \cup \partial X$ (Lemma 2.4) a subsequence of $(f(a_n))$ converges. If it converges to a point in X, then a subsequence of $(f(a_n))$ is in some compact (by properness of X) ball $\overline{B}_X(e,R)$. But then, by properness of f, a subsequence of (a_n) would be contained in some ball $B_Y(e,R')$, which would contradict $d_Y(e,a_n) \to \infty$. So some subsequence of $(f(a_n))$ converges to (that is, *represents* — see Lemma 2.4) a point in ∂X.

To establish continuity, suppose $p, q \in \partial X$. By definition of the visual metrics $d_{\partial X}$ and $d_{\partial Y}$, there exist constants $r, s > 1$ and $k, l > 0$ such that

$$d_{\partial X}(\hat{f}(p), \hat{f}(q)) \leq k r^{-(\hat{f}(p)\cdot\hat{f}(q))^X_{f(e)}} \tag{3}$$

and

$$d_{\partial Y}(p,q) \geq l s^{-(p\cdot q)_e^Y}. \tag{4}$$

Since $(p\cdot q)_e^Y = \sup\liminf_{m,n\to\infty}(p_m\cdot q_n)_e^Y$, there exist sequences (p_n),(q_m) in Y representing p,q respectively, with $\liminf_{n,m\to\infty}(p_n\cdot q_m)_e^Y \geq (p\cdot q)_e^Y - 1$. So $(p_n\cdot q_m)_e^Y \geq (p,q)_e^Y - 2$ for all sufficiently large n,m. By definition of M we have $(f(p_n)\cdot f(q_m))_{f(e)}^X \geq M((p\cdot q)_e^Y - 2)$ for such m,n, and hence

$$(\hat f(p)\cdot \hat f(q))_{f(e)}^X \geq M((p\cdot q)_e^Y - 2). \tag{5}$$

Combining (3) and (5), we have

$$d_{\partial X}(\hat f(p),\hat f(q)) \leq k r^{-M\left((p\cdot q)_e^Y-2\right)}. \tag{6}$$

So, by (4), if we make $d_{\partial Y}(p,q)$ sufficiently small, we can make $(p\cdot q)_e^Y$ arbitrarily large, so by hypothesis make $M\left((p\cdot q)_e^Y-2\right)$ arbitrarily large, and so by (6) make $d_{\partial X}(\hat f(p),\hat f(q))$ arbitrarily small.

The equivalence of (c) and (d) comes from Lemma 2.1, which implies that there exists $C>0$ such that $M'(N)\leq M(N+C)+C$ and $M(N)\leq M'(N+C)+C$ for all N.

That (d) $\Longrightarrow$ (e) is immediate as $[z,y]_Y$ is a geodesic segment in Y lying outside $B_Y(e,N)$.

Here is a proof that (e) $\Longrightarrow$ (d) under the assumption that $\sup\{d_X(f(x),f(y)) \mid d_Y(x,y)\leq r\}<\infty$ for all $r\geq 0$. Suppose $\lambda=[h_1,h_2]_Y$. As $t:=(h_1\cdot h_2)_e^Y$ approximates $d_Y(\lambda,e)$ with error at most a constant (Lemma 2.1), it is enough to show $d_X([f(h_1),f(h_2)]_X,e)\to\infty$ as $t\to\infty$. Let $\alpha_i=[e,h_i]_Y$ for $i=1,2$. By the slim–triangles condition, $[f(h_1),f(h_2)]_X$ lies in a C–neighborhood of a piecewise–geodesic path $[f(h_1),f(\alpha_1(t))]_X\cup[f(\alpha_1(t)),f(\alpha_2(t))]_X\cup[f(\alpha_2(t)),f(h_2)]_X$ for some constant C. So it is enough to show that the distance of each of these three segments from $f(e)$ in X tends to ∞ as $t\to\infty$. This is so for $[f(h_1),f(\alpha_1(t))]_X$ and $[f(\alpha_2(t)),f(h_2)]_X$ by (e). By the thin–triangles condition (see e.g. [4, pages 408–409]), $d_Y(\alpha_1(l),\alpha_2(l))$ is at most a constant for all $0\leq l\leq t$, and so in particular $d_Y(\alpha_1(t),\alpha_2(t))$ is at most a constant. The assumption that $\sup\{d_X(f(x),f(y))\mid d_Y(x,y)\leq r\}<\infty$ for all $r\geq 0$ gives an upper bound, independent of t, on the length of $[f(\alpha_1(t)),f(\alpha_2(t))]_X$. Since the distances of the endpoints of this segment from $f(e)$ in X tend to ∞ as $t\to\infty$, so does the distance of the whole segment. □

The first part of the following lemma shows that Theorem 1.1 is not a quirk of the choice of generating sets. The second establishes the sense in which the function $\varepsilon(\delta)$ of Section 1 is an invariant for Cannon–Thurston maps. The third will allow us to reinterpret Lemma 3.1 (as Corollary 3.3) in a manner well suited to analyzing hyperbolic hydra.

Lemma 3.2. *Suppose Λ is a hyperbolic subgroup of a hyperbolic group Γ.*

(i). *Whether the Cannon–Thurston map $\partial\Lambda\to\partial\Gamma$ is well–defined does not depend on the choice of finite generating sets giving the word metrics.*

(ii). *If the Cannon–Thurston map $\imath:\partial\Lambda\to\partial\Gamma$ is well–defined for a hyperbolic subgroup Λ of a hyperbolic group Γ, the modulus of continuity for $\imath$ does not depend on the finite generating sets and the choices of visual metrics up to the following Hölder–type equivalence. If $\varepsilon(\delta)$ and $\varepsilon'(\delta)$ are the moduli of continuity of $\imath$ defined with respect to different such choices, then there are functions $f_1,f_2:(0,\infty)\to(0,\infty)$, each of the form $x\mapsto C_i x^{\alpha_i}$ for some $C_i,\alpha_i>0$, such that $\varepsilon'(\delta)\leq(f_1\circ\varepsilon\circ f_2)(\delta)$ for all $\delta>0$.*

(iii). *Suppose A and B are finite generating sets for Γ and Λ, respectively. Suppose $f: C_B(\Lambda)\to C_A(\Gamma)$ is any map between the respective Cayley graphs which restricts to the inclusion $\Lambda\hookrightarrow\Gamma$ on the vertices of $C_B(\Lambda)$ and sends edges to geodesics in $C_A(\Gamma)$. The Cannon–Thurston map $\partial\Lambda\to\partial\Gamma$, defined in terms of finite generating sets A for Γ and B for Λ, is well–defined if and only if the Cannon–Thurston map $\partial C_B(\Lambda)\to\partial C_A(\Gamma)$ is.*

Proof. Suppose $\partial_1\Gamma$, $\partial_2\Gamma$, $\partial_1\Lambda$, and $\partial_2\Lambda$ are boundaries of Γ and Λ defined with respect to different finite generating sets. The composition $\partial_1\Lambda \to \partial_2\Lambda \to \partial_2\Gamma \to \partial_1\Gamma$ of Cannon–Thurston maps is $\partial_1\Lambda \to \partial_1\Gamma$, so (i) follows since $\partial_1\Lambda \to \partial_2\Lambda$ and $\partial_2\Gamma \to \partial_1\Gamma$ are well–defined since they are induced by identity maps on Λ and Γ changing the word metrics, which are Lipschitz.

If X, Y, Z are metric spaces and $g : Y \to Z$ and $h : X \to Y$ are maps with moduli of continuity ε_g and ε_h, respectively, then it follows from the definition that $\varepsilon_{g\circ h}(\delta) \leq (\varepsilon_g \circ \varepsilon_h)(\delta)$ for all $\delta > 0$. Specializing to the case of $\partial_1\Lambda \to \partial_2\Lambda \to \partial_2\Gamma \to \partial_1\Gamma$ in the previous paragraph, (ii) follows from the fact (Lemma 2.3) that $\partial_1\Lambda \to \partial_2\Lambda$ and $\partial_2\Gamma \to \partial_1\Gamma$ are Hölder.

For (iii), note that the map identifying Γ (resp. Λ) with the vertex set of $C_A(\Gamma)$ (resp. $C_B(\Lambda)$) induces an isometry on their boundaries. □

Corollary 3.3. *Suppose Λ is a finite–rank free subgroup of a hyperbolic group Γ. Suppose A is a finite generating set for Γ and B is a free basis for Λ. The Cannon–Thurston map $\partial\Lambda \to \partial\Gamma$ is well–defined if and only if for all $M'' > 0$, there exists N such that whenever $\alpha\beta$ is a reduced word on B with $|\alpha| \geq N$, any geodesic in the Cayley graph $C_A(\Gamma)$ joining α to $\alpha\beta$ lies outside the ball of radius M'' about e.*

Proof. By Lemma 3.2(iii), the Cannon–Thurston map $\partial\Lambda \to \partial\Gamma$ is well–defined if and only if the Cannon–Thurston map $\partial C_B(\Lambda) \to \partial C_A(\Gamma)$ for f (as defined in that lemma) is. Now applying condition (e) of Lemma 3.1 to f gives the result, since reduced words correspond to geodesics in $C_B(\Lambda)$. □

4 Cannon–Thurston maps for hyperbolic hydra groups

In this section we will show that Cannon–Thurston maps for hyperbolic hydra are well–defined. Throughout we fix an integer $k \geq 1$.

The *hyperbolic hydra* Γ_k of [3] is a hyperbolic group

$$\Gamma_k = F \rtimes_\theta \mathbb{Z}$$

where $F := F(a_0, \ldots, a_k, b_1, \ldots, b_l)$, and $l \geq 1$ is a certain integer, and θ is a certain automorphism of F whose restriction to $F(b_1, \ldots, b_l)$ is an automorphism and

$$\theta(a_i) = \begin{cases} u a_1 v & i = 0, \\ a_0 & i = 1, \\ a_i a_{i-1} & 1 < i \leq k, \end{cases}$$

for certain words u and v on $b_1, \ldots, b_l$. (In fact u and v depend on k, but l and $\theta|_{F(b_1,\ldots,b_l)}$ do not.)

Let t denote a generator of the $\mathbb{Z}$–factor, so $t^{-1}a_i t = \theta(a_i)$ and $t^{-1}b_j t = \theta(b_j)$ for all i and j. For $1 \leq r \leq k$, let Λ_r be the subgroup $\langle a_0 t, \ldots, a_r t, b_1, \ldots, b_l \rangle$ of Γ_k. It is proved in [3] that Λ_k is free of rank $k + l + 1$ and is distorted so that $\mathrm{Dist}_{\Lambda_k}^{\Gamma_k} \succeq A_k$. Let $\Lambda_0 = \langle b_1, \ldots, b_l \rangle$. (*Caution:* here we have that $\Lambda_r \leq \Lambda_k \leq \Gamma_k$ in contrast to [3] where $\Lambda_r \leq \Gamma_r$.)

The hyperbolic hydra Γ_k is an elaboration of the *hydra group*

$$G_k = F(a_1, \ldots, a_k) \rtimes_\varphi \mathbb{Z}$$

from [8], where t generates the $\mathbb{Z}$–factor and φ is the automorphism of $F(a_1, \ldots, a_k)$ such that

$$\varphi(a_i) = t^{-1} a_i t = \begin{cases} a_1 & i = 1, \\ a_i a_{i-1} & 1 < i \leq k. \end{cases}$$

In [8] it is proved that G_k is CAT(0) and has a rank–k free subgroup $H_k = \langle a_1 t, \ldots, a_k t \rangle$, distorted so that $\mathrm{Dist}_{H_k}^{G_k} \simeq A_k$. While G_k is easier to work with than Γ_k, it fails to be hyperbolic. The groups are related by the map $\Phi : \Gamma_k \twoheadrightarrow G_k$ which sends $a_i \mapsto a_{\max\{1,i\}}$, $b_j \mapsto 1$ and $t \mapsto t$ for all i, j.

The *normal form* of g in G_k (resp. Γ_k) is the unique $\widehat{w}\, t^m$ such that $\widehat{w}$ is a reduced word on $a_1, \ldots, a_k$ (resp. $a_0, \ldots, a_k, b_1, \ldots, b_l$) and $g = \widehat{w}\, t^m$ in G_k (resp. Γ_k). For g in G_k (resp. Γ_k) we denote the length of a shortest word on $a_1, \ldots, a_k, t$ (resp. $a_0, \ldots, a_k, b_1, \ldots, b_l, t$) representing g in G_k (resp. Γ_k) by $\left|g\right|_{G_k}$ (resp. $\left|g\right|_{\Gamma_k}$). An H_r*–word* (resp. Λ_r*–word*) is a reduced word on $a_1 t, \ldots, a_r t$ (resp. on $a_0 t, \ldots, a_r t, b_1, \ldots, b_l$). For g in H_k (resp. Λ_k) we denote the length of a shortest H_k–word (resp. Λ_k–word) representing g by $\left|g\right|_{H_k}$ (resp. $\left|g\right|_{\Lambda_k}$).

The following lemma is a consequence of Proposition 4.8 of [3] (resp. Lemma 6.1 of [8]), which says that $\Lambda_k \cap \langle t \rangle = \{1\}$ (resp. $H_k \cap \langle t \rangle = \{1\}$).

Lemma 4.1. *Given $g \in \Gamma_k$ (resp. G_k), if there exists j such that $g t^j \in \Lambda_k$ (resp. H_k), then that j is unique.*

Corollary 4.2. *If $g_1, g_2 \in \Lambda_k$ (resp. G_k) have normal forms $\hat{w}\, t^{n_1}$ and $\hat{w}\, t^{n_2}$, respectively, then $g_1 = g_2$.*

Next we give a technical lemma comparing the location of the final a_r or a_r^{-1} in a Λ_r–word w to that in its normal form $\hat{w}\, t^m$, and moreover in $\theta^n(\widehat{w})$. (When we refer to $\theta^n(\widehat{w}) \in F$ as a *word*, we mean the reduced word on $a_0, \ldots, a_k, b_1, \ldots, b_l$ that represents it.)

Lemma 4.3. *For all integers $A, B \geq 0$ and r with $k \geq r > 1$, there exists N such that if $|n| \leq B$ and $w = u(a_r t)^{\pm 1}$ is a reduced Λ_r–word with $\ell_{\Lambda_k}(u) \geq N$, then the final $a_r^{\pm 1}$ in $\theta^n(\widehat{w})$ occurs at least A symbols in from the beginning of $\theta^n(\widehat{w})$.*

Similarly, such an N exists for the final $a_0^{\pm 1}$ or $a_1^{\pm 1}$ in $\theta^n(\widehat{w})$ when $r = 1$ and w is $u(a_0 t)^{\pm 1}$ or $u(a_1 t)^{\pm 1}$.

Proof. Assume $r > 1$.

First we explain why there is an $a_r^{\pm 1}$ in $\theta^n(\widehat{w})$ and how it relates to the $a_r^{\pm 1}$ in w. This stems from properties of H_k and G_k. If $\widehat{x}\, t^m$ is the normal form of a reduced H_r–word x, then the relative locations of the $a_r^{\pm 1}$ in $\widehat{x}$ agree with those in x. This is because when each $t^{\pm 1}$ is shuffled to the end of x, each a_i it passes is changed to $\varphi^{\mp 1}(a_i)$, and this neither introduces nor destroys any $a_r^{\pm 1}$, and it causes no cancellations between any $a_r^{\pm 1}$. (See Lemma 6.2 in [8] for more details.)

The same is true if $\widehat{x}\, t^m$ is the normal form of a reduced Λ_r–word x (when $r > 1$). After all, changing letters a_i to $\theta^{\mp 1}(a_i)$ again neither introduces nor destroys any $a_r^{\pm 1}$, and there can be no cancellations between any $a_r^{\pm 1}$, because the $\Phi : \Gamma_k \twoheadrightarrow G_k$ converts the process to one in G_k and there are no cancellations between any $a_r^{\pm 1}$ there.

Likewise, the locations of the $a_r^{\pm 1}$ in $\theta^n(\widehat{w})$, for any $n \in \mathbb{Z}$, correspond to their locations in w.

Now we will prove the lemma in the case where $w = u(a_r t)^{-1}$. Since the final letter of $\theta^{\pm 1}(a_r^{-1})$ is a_r^{-1}, the a_r^{-1} in $\theta^n(\widehat{w})$ that corresponds to the final a_r^{-1} in w is the final letter of $\theta^n(\widehat{w})$. So we are seeking to prove that $|\theta^n(\widehat{w})|_F \geq A$. Since θ is an automorphism, only finitely many $v \in F$ satisfy $|\theta^n(v)|_F < A$. By Corollary 4.2, each such v equals $\widehat{w}$ for at most one Λ–word w, so the lemma is proved by taking N sufficiently large to avoid these finitely many w.

Next we address the case where $w = u(a_r t)$. The normal forms $\widehat{w}\, t^{j+1}$ of w and $\widehat{u}\, t^j$ of u are related in that

$$w = u(a_r t) = \widehat{u}\, t^j\, a_r t = \widehat{u}\, \theta^{-j}(a_r)\, t^{j+1} = \widehat{w}\, t^{j+1}.$$

Since there is no cancellation between u and $a_r t$ in w and the first letter of $\theta^{-j}(a_r)$ is a_r, there is no cancellation between $\widehat{u}$ and $\theta^{-j}(a_r)$. Similarly, there is no cancellation between $\theta^n(\widehat{u})$ and $\theta^{n-j}(a_r)$ in $\theta^n(\widehat{w}) = \theta^n(\widehat{u}\, \theta^{-j}(a_r))$. So we are now seeking to prove that $|\theta^n(\widehat{u})|_F + 1 \geq A$, and this can be done as before. This completes the proof when $r > 1$.

The same proof works when $r = 1$, taking into account that θ interchanges a_0 and a_1 (introducing some $b_1, \ldots, b_l$ in the process). This is only a superficial complication: the locations of the $a_0^{\pm 1}$ and $a_1^{\pm 1}$ in $\theta^n(\widehat{w})$, for any $n \in \mathbb{Z}$, correspond to their locations in w, but some $a_0^{\pm 1}$ may correspond to $a_1^{\pm 1}$ and vice versa. $\square$

Let $C(F)$ and $C(\Gamma_k)$ denote the Cayley graphs of F and Γ_k with respect to $a_0, \ldots, a_k, b_1, \ldots, b_l$ and $a_0, \ldots, a_k, b_1, \ldots, b_l, t$. Let $B_F(e, R)$ denote the open ball of radius R about e in $C(F)$. Write $[x, y]_F$ or $[x, y]_{\Gamma_k}$ for a geodesic between x and y in $C(F)$ or $C(\Gamma_k)$. In the case of $C(F)$, which is a tree, geodesics between any given pair of points are unique. Let d_F and d_{Γ_k} be the associated metrics.

The *shadow* of the suffix β of a reduced Λ_k–word $\alpha\beta$ is the set of all geodesic segments $[\widehat{\alpha\beta_i}, \widehat{\alpha\beta_{i+1}}]_F$ where β_i denotes the length–i prefix of β and $0 \le i < |\beta|_{\Lambda_k}$.

Lemma 4.4. *For all $K > 0$, there exist integers $C, R > 0$ such that if $\alpha\beta$ is a reduced Λ_k–word and $|\alpha|_{\Lambda_k} \ge C$ and the shadow of β is outside $B_F(e, R)$, then every $[\alpha, \alpha\beta]_{\Gamma_k}$ satisfies $d_{\Gamma_k}([\alpha, \alpha\beta]_{\Gamma_k}, e) \ge K$.*

Proof. As Γ_k is hyperbolic, there is some $\delta > 0$ such that every geodesic triangle in $C(\Gamma_k)$ is δ–slim.

The geodesic segment $[\alpha, \alpha\beta]_{\Gamma_k}$ is in a 2δ–neighborhood of the piecewise–geodesic path $[\alpha, \widehat{\alpha}]_{\Gamma_k} \cup [\widehat{\alpha}, \widehat{\alpha\beta}]_{\Gamma_k} \cup [\widehat{\alpha\beta}, \alpha\beta]_{\Gamma_k}$ in $C(\Gamma_k)$. So it suffices to have $[\alpha, \widehat{\alpha}]_{\Gamma_k}$, $[\widehat{\alpha}, \widehat{\alpha\beta}]_{\Gamma_k}$, and $[\widehat{\alpha\beta}, \alpha\beta]_{\Gamma_k}$ at least $K + 2\delta$ from e.

To ensure $[\alpha, \widehat{\alpha}]_{\Gamma_k}$ is at least $K + 2\delta$ from e, we need it to avoid finitely many elements of Γ_k, say $g_1, \ldots, g_m$. Since $\alpha = \widehat{\alpha} t^n$ for some n, there is a unique geodesic $[\alpha, \widehat{\alpha}]$ joining α to $\widehat{\alpha}$ in $C(\Gamma_k)$ and it is a succession of edges all labelled t. So if g_i is on $[\alpha, \widehat{\alpha}]$ then $g_i t^{j_i} = \alpha \in \Lambda_k$ for some j_i. But then by Lemma 4.1, it suffices for α not to be one of at most m elements of Λ_k. So it suffices to ensure $|\alpha|_{\Lambda_k}$ is sufficiently long.

Since $|\alpha\beta|_{\Lambda_k} \ge |\alpha|_{\Lambda_k}$, the geodesic $[\widehat{\alpha\beta}, \alpha\beta]_{\Gamma_k}$ is then also at least $K + 2\delta$ from e.

Finally we consider $[\widehat{\alpha}, \widehat{\alpha\beta}]_{\Gamma_k}$. As the shadow of β stays outside $B_F(e, R)$, so does $[\widehat{\alpha}, \widehat{\alpha\beta}]_F$, which is a subset of the shadow since $C(F)$ is a tree. Mitra shows in his proof of Theorem 4.3 in [16] that if $1 \to H \to G \to G/H \to 1$ is a short exact sequence of finitely generated groups and H and G are infinite and hyperbolic, then condition (b) of our Lemma 3.1 applies (with X and Y being Cayley graphs of G and H). Applying this to $G = \Gamma_k$ and $H = F$, we learn that choosing R large enough makes $[\widehat{\alpha}, \widehat{\alpha\beta}]_{\Gamma_k}$ arbitrarily far from e. $\square$

We are now ready to use Corollary 3.3 to show the Cannon–Thurston map $\partial\Lambda_k \to \partial\Gamma_k$ is well–defined.

Proof of Theorem 1.1. Fix some $\delta > 0$ so that all geodesic triangles in $C(\Gamma_k)$ are δ–slim.

For integers $A, B \ge 0$ and r with $k \ge r \ge 1$, let $N(r, A, B)$ be the least integer N as per Lemma 4.3. Given an integer $R > 0$, recursively define a sequence $N_k(R), N_{k-1}(R), \ldots, N_1(R)$ of positive integers by:

$$N_k(R) := N(k, R, 0),$$

and for $r = k-1, k-2, \ldots, 1$

$$N_r(R) := N\left(r, R + \max\left\{ |\widehat{w}|_{F_k} \,\middle|\, \Lambda_k\text{–words } w \text{ with } |w|_{\Lambda_k} = \sum_{j=r+1}^{k} N_j(R) \right\}, \sum_{j=r+1}^{k} N_j(R)\right).$$

Suppose $M'' > 0$ is given. Let $K = M'' + (2\delta + 1)k$. Let $R, C > 0$ be obtained from K as per Lemma 4.4. Recall the map Φ from the hyperbolic hydra group Γ_k to the hydra group G_k defined by $a_i \mapsto a_{\max\{1,i\}}$, $b_j \mapsto 1$ and $t \mapsto t$ for all i, j. Let L be the maximum of $|\theta^n(s)|_F$ ranging over all $s \in \{b_1, \ldots, b_l\}$ and all $n \in \mathbb{Z}$ for which there exists $\widehat{u} \in F$ such that $\widehat{u} t^n \in \Lambda_k$ and the reduced word representing $\Phi(\widehat{u})$ in $F(a_1, \ldots, a_k)$ has length less than R. (There are only finitely many such n since $\widehat{u} t^n \in \Lambda_k$ implies $\Phi(\widehat{u} t^n) = \Phi(\widehat{u}) t^n \in H_k$, and for any $x \in F(a_1, \ldots, a_k)$, there is at most one $m \in \mathbb{Z}$ such that $x t^m \in H_k$ by Lemma 4.1.) Choose $C' > C$ so that every Λ_k–word u of length $|u|_{\Lambda_k} \ge C'$ has free–by–cyclic normal form $\widehat{u} t^n$ with $\left|\widehat{u}\right|_F \ge R + (L/2)$.

Suppose $\alpha\beta$ is a reduced Λ_k–word such that $|\alpha|_{\Lambda_k} \geq N := C' + \sum_{r=1}^{k} N_r(R)$. Express α as $w_k \cdots w_0$ where $|w_r|_{\Lambda_k} = N_r(R)$ for $1 \leq r \leq k$. In particular, $|\alpha|_{\Lambda_k} \geq |w_0|_{\Lambda_k} \geq C' \geq C$. Let β_r denote the longest prefix of β in Λ_r and let γ_r denote (any) geodesic $[\alpha, \alpha\beta_r]_{\Gamma_k}$. In particular, γ_k is (any) geodesic $[\alpha, \alpha\beta]_{\Gamma_k}$. We will show that γ_k lies at least a distance M'' from e in $C(\Gamma_k)$. Corollary 3.3 will then complete the proof.

Suppose, for a contradiction, that $d_{\Gamma_k}(\gamma_k, e) < M''$. Suppose $\widehat{\alpha} t^n$ is the free–by–cyclic normal form of α in Γ_k.

We show in this paragraph that the shadow of the suffix β_0 of $\alpha\beta_0$ is outside of $B_F(e,R)$. The endpoints of the geodesic segments in F comprising the shadow of β_0 are all of the form $\widehat{\alpha}\theta^{-n}(x)$ for various $x \in F(b_1, \ldots, b_l)$. There are two cases to consider: the length of the reduced word in $F(a_1, \ldots, a_k)$ representing $\Phi(\widehat{\alpha})$ is at least R and is less than R. In the former case, because $\widehat{\alpha}$ contains at least R letters $a_i^{\pm 1}$ $(0 \leq i \leq k)$, the closest approach of any such geodesic to e (the Gromov product of its endpoints) is at least R. In the latter case, L is an upper bound for the length of the constituent geodesics in the shadow of β_0, and so, by definition of C', the shadow of β_0 stays outside of $B_F(e,R)$. In either case, the shadow stays outside of $B_F(e,R)$.

On the other hand, the following claim, in the case $r = 0$, shows that

$$d_{\Gamma_k}([\alpha, \alpha\beta_0]_{\Gamma_k}, e) \ < \ M'' + (2\delta + 1)k,$$

which equals K, so Lemma 4.4 implies the shadow of β_0 dips within $B_F(e,R)$. This contradiction will prove the theorem.

Claim. For $r = k, k-1, \ldots, 1, 0$,

(i_r). $d_{\Gamma_k}(\gamma_r, e) < M'' + (2\delta + 1)(k - r)$, and

(ii_r). $w_r w_{r-1} \cdots w_0 \in \Lambda_r$.

We prove this claim using simultaneous downward induction on r.

The base case $r = k$ is straightforward: $\gamma_k = [\alpha, \alpha\beta]_{\Gamma_k}$ and $w_k \cdots w_0 = \alpha \in \Lambda_k$ by definition, and $d_{\Gamma_k}(\gamma_k, e) < M''$ by hypothesis.

Now we prove that (i_{r+1}) and (ii_{r+1}) implies (i_r) and (ii_r) for $r = k-1, \ldots, 0$, except we will not give the case $r = 0$ explicitly; it is similar to the following $r > 0$ case but has the superficial complication that θ interchanges $a_0^{\pm 1}$ and $a_1^{\pm 1}$ (as well as introducing some letters $b_j^{\pm 1}$), which makes the notation more awkward.

We will make repeated use of the following lemma.

Lemma 4.5. *Suppose $w_k, \ldots, w_{r+1}$ are as defined earlier, and $w_k \cdots w_{r+1} x (a_{r+1} t)^{\pm 1} y z$ is a reduced Λ_k–word in which the subwords x, y and z are Λ_{r+1}–words. If γ is any geodesic in $C(\Gamma_k)$ from $w_k \cdots w_{r+1} x (a_{r+1} t)^{\pm 1} y$ to $w_k \cdots w_{r+1} x (a_{r+1} t)^{\pm 1} y z$, then $d_{\Gamma_k}(\gamma, e) \geq K$.*

To prove this lemma we consider the free–by–cyclic normal form $\widehat{w_k \cdots w_{r+2}}\, t^{n_r}$ for $w_k \cdots w_{r+2}$ in Γ_k. Now, $|w_k \cdots w_{r+2}|_{\Lambda_k} = \sum_{j=r+2}^{k} N_j(R)$ by construction, so $|n_r| \leq \sum_{j=r+2}^{k} N_j(R)$. So, by definition,

$$N_{r+1}(R) \ \geq \ N(r+1, R + |\widehat{w_k \cdots w_{r+2}}|_F, |n_r|).$$

By hypothesis, $w_{r+1} x (a_{r+1} t)^{\pm 1} \in \Lambda_{r+1}$, and $|w_{r+1}| = N_{r+1}(R)$ by construction, so the final $a_{r+1}^{\pm 1}$ in $\theta^{-n_r}(\widehat{x(a_{r+1}t)^{\pm 1}})$ occurs at least $R + |\widehat{w_k \cdots w_{r+2}}|_F$ symbols in from the start of the word. Therefore, the final $a_{r+1}^{\pm 1}$ in $\widehat{w_k \cdots w_{r+2}}\ \theta^{-n_r}(\widehat{x(a_{r+1}t)^{\pm 1}})$ occurs at least $(R + |\widehat{w_k \cdots w_{r+2}}|_F) - |\widehat{w_k \cdots w_{r+2}}|_F = R$ symbols in. Since y and z are Λ_{r+1}–words and $w_k \cdots w_{r+1} x (a_{r+1} t)^{\pm 1} y z$ is reduced, this implies the shadow of z cannot dip inside $B_F(e,R)$. So, by Lemma 4.4, $d_{\Gamma_k}(\gamma, e) \geq K$, completing the proof of the lemma.

Returning to the proof of the claim, we will consider two cases: $\beta_r = \beta_{r+1}$ and $\beta_r \neq \beta_{r+1}$. In the former case, we may assume $\gamma_r = \gamma_{r+1}$. So (i_r) follows immediately from (i_{r+1}). Since (i_r) implies $d_{\Gamma_k}(\gamma_r, e) < K$, (ii_{r+1}) and Lemma 4.5 with $z = \beta_r$ shows that $w_r \cdots w_0$ cannot be expressed as $x(a_{r+1}t)^{\pm 1} y$. So $w_r \cdots w_0 \in \Lambda_r$ and we have (ii_r).

Next, assume $\beta_r \neq \beta_{r+1}$. The (reduced) word β_{r+1} can be expressed as $\beta_r(a_{r+1}t)^{\pm 1}\beta_r'$ for some $\beta_r' \in \Lambda_{r+1}$. Let ρ_r be the geodesic segment in $C(\Gamma_k)$ labelled $(a_{r+1}t)^{\pm 1}$ connecting $\alpha\beta_r$ to $\alpha\beta_r(a_{r+1}t)^{\pm 1}$. Let $\gamma_r' = [\alpha\beta_r(a_{r+1}t)^{\pm 1}, \alpha\beta_{r+1}]_{\Gamma_k}$. Then by Lemma 4.5 with $x = w_r \cdots w_0 \beta_r$, y the empty word, and $z = \beta_{r+1}'$,

$$d_{\Gamma_k}(\gamma_r', e) \ \geq \ K. \tag{7}$$

If $(a_{r+1}t)^{\pm 1}$ occurs in $w_r \cdots w_0$, then Lemma 4.5 with $x(a_{r+1}t)^{\pm 1}y = w_r \cdots w_0$ and $z = \beta_r$ shows that

$$d_{\Gamma_k}(\gamma_r, e) \ \geq \ K. \tag{8}$$

By the slim–triangles condition for $C(\Gamma_k)$, γ_{r+1} is contained in the 2δ–neighborhood of $\gamma_r \cup \rho_r \cup \gamma_r'$ and hence in the $(2\delta+1)$–neighborhood of $\gamma_r \cup \gamma_r'$. So

$$\begin{aligned} \min\{d(e,\gamma_r), d(e,\gamma_r')\} \ &\leq \ d(e,\gamma_{r+1}) + (2\delta+1) \\ &< \ M'' + (2\delta+1)(k-(r+1)) + (2\delta+1) \\ &= \ M'' + (2\delta+1)(k-r), \end{aligned}$$

the second inequality coming from (i_{r+1}).

But by (7), $d(e,\gamma_r') \geq K \geq M'' + (2\delta+1)(k-r)$, so $\min\{d(e,\gamma_r), d(e,\gamma_r')\} = d(e,\gamma_r)$ and (i_r) follows.

Moreover, (8) cannot be true since it contradicts (i_r), so $(a_{r+1}t)^{\pm 1}$ does not occur in $w_r \cdots w_0$ and (ii_r) follows. This completes the induction step of the claim, so the theorem is proved. ☐

5 Wildness of Cannon–Thurston maps

Proof of Theorem 1.2. We have that Γ and Λ are (δ_Γ)– and (δ_Λ)–hyperbolic, respectively, for some $\delta_\Gamma, \delta_\Lambda > 0$. Let $\imath : \Lambda \to \Gamma$ denote the inclusion map and $\hat{\imath} : \partial\Lambda \to \partial\Gamma$ denote the Cannon–Thurston map.

Since Λ is non–elementary, $|\hat{\imath}(\partial\Lambda)| = \infty$ by [11, Thm. 12.2(1)]. We may thus choose $p_1, p_2, p_3 \in \partial\Lambda$ with $\hat{\imath}p_1, \hat{\imath}p_2, \hat{\imath}p_3 \in \partial\Gamma$ distinct. Let $C = 2\delta_\Lambda + \max\{(p_i \cdot p_j)_e^\Lambda \mid 1 \leq i < j \leq 3\}$.

By definition of the distortion function (see Section 1), we can take a sequence $h_n \in \Lambda$ with $d_\Gamma(e, h_n) \leq n$ and $d_\Lambda(e, h_n) = \mathrm{Dist}_\Lambda^\Gamma(n)$. By Lemma 2.2(1)

$$(p_i \cdot p_j)_e^\Lambda \ = \ (h_n p_i \cdot h_n p_j)_{h_n}^\Lambda \ \geq \ \min\{(e \cdot h_n p_i)_{h_n}^\Lambda, (e \cdot h_n p_j)_{h_n}^\Lambda\} - 2\delta_\Lambda,$$

so we can choose $i = i(n)$ and $j = j(n) \in \{1,2,3\}$ with $i \neq j$ such that

$$(e \cdot h_n p_i)_{h_n}^\Lambda, (e \cdot h_n p_j)_{h_n}^\Lambda \ \leq \ C.$$

Combined with Lemma 2.2(2) this gives that for $k = i, j$,

$$(h_n \cdot h_n p_k)_e^\Lambda \ \geq \ d_\Lambda(e, h_n) - (e \cdot h_n p_k)_{h_n}^\Lambda - \delta_\Lambda \ \geq \ \mathrm{Dist}_\Lambda^\Gamma(n) - C - \delta_\Lambda.$$

So, by Lemma 2.2(1),

$$(h_n p_i \cdot h_n p_j)_e^\Lambda \ \geq \ \min\{(h_n \cdot h_n p_i)_e^\Lambda, (h_n \cdot h_n p_j)_e^\Lambda\} - 2\delta_\Lambda \ \geq \ \mathrm{Dist}_\Lambda^\Gamma(n) - C - 3\delta_\Lambda.$$

Writing $\beta := k_2 r^{C+3\delta_\Lambda}$, where k_2 is as per (2) in Section 2 applied to $\partial\Lambda$, we get

$$d_{\partial\Lambda}(h_n p_i, h_n p_j) \ \leq \ k_2 r^{-(h_n p_i \cdot h_n p_j)_e^\Lambda} \ \leq \ \frac{\beta}{r^{\mathrm{Dist}_\Lambda^\Gamma(n)}}. \tag{9}$$

On the other hand, using Lemma 2.2(3) for the first inequality,

$$\begin{aligned}(\hat{\imath}(h_n p_i)\cdot\hat{\imath}(h_n p_j))^{\Gamma}_e &= (h_n\hat{\imath}p_i\cdot h_n\hat{\imath}p_j)^{\Gamma}_e\\ &\leq d_\Gamma(e,[h_n\hat{\imath}p_i,h_n\hat{\imath}p_j]_\Gamma)+8\delta_\Gamma\\ &\leq d_\Gamma(e,h_n)+d_\Gamma(h_n,[h_n\hat{\imath}p_i,h_n\hat{\imath}p_j]_\Gamma)+8\delta_\Gamma\\ &\leq n+d_\Gamma(e,[\hat{\imath}p_i,\hat{\imath}p_j]_\Gamma)+8\delta_\Gamma.\end{aligned}$$

Writing $\alpha := k_1/s^{8\delta_\Gamma+\max\{d_\Gamma(e,[\hat{\imath}p_i,\hat{\imath}p_j]_\Gamma)|1\leq i<j\leq 3\}}$, where k_1 is as per (2) in Section 2 applied to $\partial\Gamma$, we then get

$$d_{\partial\Gamma}(\hat{\imath}(h_n p_i),\hat{\imath}(h_n p_j)) \geq k_1 s^{-(\hat{\imath}(h_n p_i),\hat{\imath}(h_n p_j))^{\Gamma}_e} \geq \frac{\alpha}{s^n}. \tag{10}$$

Combining (9) and (10) yields the inequality claimed in Theorem 1.2. ☐

Our final proposition provides precise details of what we mean by Corollary 1.3 via $\delta = 1/n$. The inequality we obtain is not as clean as we might like, but for $k=3$ (when it is a tower of powers of 2) and beyond, the Ackermann function A_k is by far the dominant force on the lefthand side, and can be considered to absorb the log and the exponential.

Proposition 5.1. *For all $k\geq 2$, the modulus of continuity $\varepsilon(\delta)$ of the Cannon–Thurston map $\partial\Lambda_k\to\partial\Gamma_k$ for hyperbolic hydra has the property that there exist $C_0, C_1>0$ and $C_2>1$ such that for all $\eta\in(0,C_0)$,*

$$\varepsilon\left(\frac{1}{C_2^{A_k(\lfloor C_1\log(1/\eta)\rfloor)}}\right) \geq \eta.$$

Proof. For convenience, extend the domains of the functions $A_k:\mathbb{N}\to\mathbb{N}$ and $\mathrm{Dist}^{\Gamma_k}_{\Lambda_k}:\mathbb{N}\to\mathbb{N}$ to $[1,\infty)$ by declaring the functions to be constant on the half–open intervals $[n,n+1)$.

From Theorem 1.2 we have

$$\varepsilon\left(\frac{\beta}{r^{\mathrm{Dist}^{\Gamma_k}_{\Lambda_k}(n)}}\right) \geq \frac{\alpha}{s^n}$$

for all real $n\geq 0$ and some constants $\alpha,\beta>0$ and $r,s>1$. Thus for all $\eta\leq\alpha$,

$$\varepsilon\left(\frac{1}{\exp\left(\log(r)\mathrm{Dist}^{\Gamma_k}_{\Lambda_k}(\log_s(\alpha/\eta))-\log(\beta)\right)}\right) \geq \eta. \tag{11}$$

As $\mathrm{Dist}^{\Gamma_k}_{\Lambda_k}\succeq A_k$ by [3], there exists $C>0$ such that $A_k(n)\leq C\mathrm{Dist}^{\Gamma_k}_{\Lambda_k}(Cn+C)+Cn+C$ for all real n, and therefore

$$\mathrm{Dist}^{\Gamma_k}_{\Lambda_k}(N) \geq \frac{1}{C}A_k\left(\frac{N-C}{C}\right)-\frac{N}{C}$$

for $N\geq 2C$. So

$$\log(r)\mathrm{Dist}^{\Gamma_k}_{\Lambda_k}(\log_s(\alpha/\eta))-\log(\beta) \geq K_1A_k(K_2\log(1/\eta)-K_3)-K_4\log(1/\eta)-K_5$$

for all $\eta\in(0,K_0)$, for suitable constants $K_0,\ldots,K_5>0$. By shrinking K_0 if necessary, we can make $\log(1/\eta)$ arbitrarily large, so that we may absorb the constant K_3 into K_2. Moreover, A_k grows faster than a linear function as $k\geq 2$, so (by further shrinking K_0) the constants K_4 and K_5 can be absorbed into K_1. Thus

$$\log(r)\mathrm{Dist}^{\Gamma_k}_{\Lambda_k}(\log_s(\alpha/\eta))-\log(\beta) \geq K_1A_k(K_2\log(1/\eta)). \tag{12}$$

So combining (11) and (12) and setting $C_0=K_0$, $C_1=K_2$, and $C_2=e^{K_1}$, we have the result claimed in the proposition. ☐

References

[1] O. Baker and T. R. Riley. Cannon–Thurston maps do not always exist. arXiv:1206.0505, 2012.

[2] M. Bestvina and G. Mess. The boundary of negatively curved groups. *J. Amer. Math. Soc.*, 4(3):469–481, 1991.

[3] N. Brady, W. Dison, and T. R. Riley. Hyperbolic hydra. arxiv:1105.1535, 2011.

[4] M. R. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1999.

[5] P. Brinkmann. Splittings of mapping tori of free group automorphisms. *Geometriae Dedicata*, 93:191–203, 2002.

[6] K. S. Brown. *Cohomology of groups*, volume 87 of *Graduate Texts in Mathematics*. Springer–Verlag, New York, 1994. Corrected reprint of the 1982 original.

[7] J. W. Cannon and W. P. Thurston. Group invariant Peano curves. *Geom. Topol.*, 11:1315–1355, 2007.

[8] W. Dison and T. R. Riley. Hydra groups. arXiv:1002.1945.

[9] E. Ghys and P. de la Harpe, editors. *Sur les Groups Hyperbolic d'après Mikhael Gromov*, volume 83 of *Progress in Mathematics*. Birkhäuser, 1990.

[10] M. Gromov. Hyperbolic groups. In S. M. Gersten, editor, *Essays in group theory*, volume 8 of *MSRI publications*, pages 75–263. Springer–Verlag, 1987.

[11] I. Kapovich and N. Benakli. Boundaries of hyperbolic groups. In *Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001)*, volume 296 of *Contemp. Math.*, pages 39–93. Amer. Math. Soc., Providence, RI, 2002.

[12] I. Kapovich and M. Lustig. On the fibers of the Cannon–Thurston map for free-by-cyclic groups. arXiv:1207.3494, 2012.

[13] M. Kapovich and B. Kleiner. Hyperbolic groups with low–dimensional boundary. *Ann. Scient. Éc. Norm. Sup.*, pages 647–669, 2000.

[14] H. Lebesgue. Sur les intégrales singulières. *Ann. Fac. Sci. Toulouse Sci. Math. Sci. Phys. (3)*, 1:25–117, 1909.

[15] Y. Matsuda and S. Oguni. On Cannon–Thurston maps for relatively hyperbolic groups. arXiv:1206.5868, 2012.

[16] M. Mitra. Cannon–Thurston maps for hyperbolic group extensions. *Topology*, 37(3):527–538, 1998.

[17] M. Mitra. Cannon–Thurston maps for trees of hyperbolic metric spaces. *J. Differential Geometry*, 48:135–164, 1998.

[18] H. Miyachi. Moduli of continuity of Cannon–Thurston maps. In *Spaces of Kleinian groups*, volume 329 of *London Math. Soc. Lecture Note Ser.*, pages 121–149. Cambridge Univ. Press, Cambridge, 2006.

[19] H. Short *et al.*. Notes on word hyperbolic groups. In E. Ghys, A. Haefliger, and A. Verjovsky, editors, *Group Theory from a Geometrical Viewpoint (Trieste, 1990)*, volume 8 of *MSRI publications*, pages 3–63. World Scientific Publishing, River Edge, N.J., 1991. www.cmi.univ-mrs.fr/~hamish/Papers/MSRInotes2004.pdf.